ALID: Scalable Dominant Cluster Detection

Lingyang Chu
Institute of Computing Technology, CAS
Beijing, China
lingyang.chu@vipl.ict.ac.cn

Shuhui Wang
Institute of Computing Technology, CAS
Beijing, China
wangshuhui@ict.ac.cn

Siyuan Liu
Heinz College
Carnegie Mellon University
Pittsburgh, USA
siyuan@cmu.edu

Qingming Huang
University of Chinese Academy of Sciences
Beijing, China
qmhuang@jdl.ac.cn

Jian Pei
School of Computing Science
Simon Fraser University
Vancouver, Canada
jpei@cs.sfu.ca

arXiv:1411.0064v1 [cs.DB] 1 Nov 2014

ABSTRACT

Detecting dominant clusters is important in many analytic applications. The state-of-the-art methods find dense subgraphs on the affinity graph as the dominant clusters. However, the time and space complexity of those methods are dominated by the construction of the affinity graph, which is quadratic with respect to the number of data points, and thus impractical on large data sets. To tackle the challenge, in this paper, we apply *Evolutionary Game Theory* (EGT) and develop a scalable algorithm, Approximate Localized Infection Immunization Dynamics (ALID). The major idea is to perform Localized Infection Immunization Dynamics (LID) to find dense subgraph within local range of the affinity graph. LID is further scaled up with guaranteed high efficiency and detection quality by an estimated Region of Interest (ROI) and a carefully designed Candidate Infective Vertex Search method (CIVS). ALID only constructs small local affinity graphs and has a time complexity of $\mathcal{O}(C(a^* + \delta)n)$ and a space complexity of $\mathcal{O}(a^*(a^* + \delta))$, where a^* is the size of the largest dominant cluster and $C \ll n$ and $\delta \ll n$ are small constants. We demonstrate by extensive experiments on both synthetic data and real world data that ALID achieves state-of-the-art detection quality with much lower time and space cost on single machine. We also demonstrate the encouraging parallelization performance of ALID by implementing the Parallel ALID (PALID) on *Apache Spark*. PALID processes 50 million SIFT data points in 2.29 hours, achieving a speedup ratio of 7.51 with 8 executors.

1. INTRODUCTION

A dominant cluster is a group of highly similar objects that possesses maximal inner group coherence [18, 26]. On a massive data set, more often than not the dominant clusters carry useful information and convey important knowledge. For example, in a big collection of news data (e.g., official news, RSS-feeds and tweet-streams), the dominant clusters may indicate potential real world hot events [4, 29]; in a large repository of interpersonal communication data (e.g., emails and social networks), the dominant clusters may reveal stable social hubs [36]. Therefore, efficiently and effectively detecting dominant clusters from massive data sets has become an important task in data analytics.

In real applications, dominant cluster detection often faces two challenges. First, the number of dominant clusters is often unknown. Second, large data sets are often noisy. Unknown number of meaningful dominant clusters are often hidden deeply in an overwhelming amount of background noise [29, 36]. For instance, numerous news items about almost every aspect of our daily life are added to the Web everyday. Most of the news items are interesting to small groups of people, and hardly attract sufficient social attention or become a dominant cluster of a hot event. As another example, billions of spam messages are sent everyday. Finding meaningful email threads and activities is a typical application of dominant cluster detection, and is challenging mainly due to the large amount of spam messages as noise [18].

Traditional partitioning-based clustering methods like k-means [5, 21, 30] and spectral clustering [15, 17, 24, 34] are often used in cluster detection. However, such methods are not robust in processing noisy data with an unknown number of dominant clusters [26]. First, these methods typically require a pre-defined number of (dominant) clusters, without prior knowledge on the true number of (dominant) clusters, an improper number of clusters may lead to low detection accuracy. Second, each data item, including both members of dominant clusters and noise data items, are forced to be assigned to a certain cluster, which inevitably leads to degenerated detection accuracy and subtracted cluster coherence under high background noise.

The affinity-based methods [18, 26, 28], which detect dominant clusters by finding dense subgraphs on an affinity graph, are effective in detecting an unknown number of dominant clusters from noisy background. Since the data objects in a dominant cluster are very similar to each other, they naturally form a highly cohesive dense subgraph on

the affinity graph. Such high cohesiveness is proven to be a stable characteristic to accurately identify dominant clusters from noisy background without knowing the exact number of clusters [18, 28].

Nevertheless, efficiency and scalability are the bottlenecks of the affinity-based methods, since the complexity of constructing the affinity matrix from n data items is $\mathcal{O}(n^2)$ in both time and space. Even though the computational costs can be saved by forcing the affinity graph sparse [8], the enforced sparsity breaks the intrinsic cohesiveness of dense subgraphs and consequently affects the detection quality of dominant clusters.

To tackle the challenge, in this paper, we propose Approximate Localized Infection Immunization Dynamics (ALID), a dominant cluster detection approach that achieves high scalability and retains high detection quality. The key idea is to avoid constructing the global complete affinity graph. Instead, ALID finds a dense subgraph within an accurately estimated local Region of Interest (ROI). The ROI is guaranteed by the *law of triangle inequality* to completely cover a dominant cluster, and thus fully preserves the high intrinsic cohesiveness of the corresponding dense subgraph and ensures the detection quality. Moreover, since dominant clusters generally exist in small local ranges, ALID only searches a small local affinity subgraph within the ROI. Therefore, ALID only constructs small local affinity graphs and largely avoids the expensive construction of the global affinity graph. Consequently, the original $\mathcal{O}(n^2)$ time and space complexity of the affinity graph construction is significantly reduced to $\mathcal{O}(C(a^*+\delta)n)$ in time and $\mathcal{O}(a^*(a^*+\delta))$ in space, where a^* is the size of the largest dominant cluster and $C \ll n$, $\delta \ll n$ are small constants.

We make the following major contributions. First, we propose LID to detect dense subgraphs on the local affinity graph within a ROI. LID localizes the Infection Immunization Dynamics [28] to efficiently seek dense subgraphs on a small local affinity graph. It only computes a few columns of the local affinity matrix to detect a dense subgraph without sacrificing detection quality. This significantly reduces the time and space complexity.

Second, we estimate a Region of Interest (ROI) and propose a novel Candidate Infective Vertex Search method (CIVS) to significantly improve the scalability of LID and ensure high detection quality. The estimated ROI is guaranteed to accurately identify the local range of the "*true*" dense subgraph (i.e., dominant cluster), which ensures the detection quality of ALID. The CIVS method is proposed to quickly retrieve the data items within the ROI, where all data items are efficiently indexed by Locality Sensitive Hashing (LSH) [12]. Demonstrated by extensive experiments on synthetic data and real world data, ALID achieves substantially better scalability than the other affinity-based methods without sacrificing the detection quality.

Third, we carefully design a parallelized solution on top of the MapReduce framework [13] to further improve the scalability of ALID. The promising parallelization performance of Parallel ALID (PALID) is demonstrated by the experiments on *Apache Spark* (http://spark.apache.org/). PALID can efficiently process 50 million SIFT data [22] in 2.29 hours and achieve a speedup ratio of 7.51 with 8 executors.

For the rest of the paper, we review related work in Section 2. We revisit the problem of dense subgraph finding in Section 3. We present the ALID method in Section 4. Section 5 reports the experimental results. Section 6 concludes the paper. We also provide a full version of this work in [10] with an appendix of proof materials and look-up tables.

2. RELATED WORK

For dominant cluster detection, the affinity-based methods [18, 26, 28] that find dense subgraphs on affinity graphs are more resistant against background noise than the canonical partitioning-based methods like k-means [5, 21] and spectral clustering [15, 17, 34]. The dense subgraph seeking problem is well investigated in literature [3, 4, 33]. Motzkin *et al.* [23] proved that seeking dense subgraphs on an unweighted graph can be formulated as a quadratic optimization problem on the simplex. This method was extended to weighted graphs by the dominant set method (DS) [26], which solves a standard quadratic optimization problem (StQP) by replicator dynamics (RD) [35].

Bulò *et al.* [28] showed that, given the full affinity matrix of an affinity graph with n vertices, the time complexity for each RD iteration is $\mathcal{O}(n^2)$, which hinders its application on large data sets. Thus, Bulò *et al.* [28] proposed the infection immunization dynamics (IID) [28] to solve the StQP problem in $\mathcal{O}(n)$ time and space. However, the overall time and space complexity of IID is still $\mathcal{O}(n^2)$, since each iteration of IID needs the full affinity matrix, which costs quadratic time and space to compute and store.

Since most dense subgraphs exist in local ranges of an affinity graph, running RD on the entire graph is inefficient [18, 19]. Therefore, Liu *et al.* [18] proposed the shrinking and expansion algorithm (SEA) to effectively prevent unnecessary time and space cost by restricting all RD iterations on small subgraphs. Both the time and space complexities of SEA are linear with respect to the number of graph edges [18]. The scalability of SEA is sensitive to the sparse degree of the affinity graph.

Affinity propagation (AP) [16] is another noise resistant solution to detect unknown number of dominant clusters. It finds the dominant clusters by passing real valued messages along graph edges, which is very time consuming when there are many vertices and edges.

Mean shift [11] differs from the affinity-based methods by directly seeking clusters in the feature space. It assumes that the discrete data items are sampled from a pre-defined density distribution and detect clusters by iteratively seeking the maxima of the density distribution. However, the detection quality of mean shift is sensitive to the type and bandwidth of the pre-defined density distribution.

The above affinity-based methods are able to achieve high detection quality when the affinity matrix is already materialized. However, the scalability of those methods with respect to large data sets is limited by the $\mathcal{O}(n^2)$ time and space complexities of the affinity matrix computation. To the best of our knowledge, ALID developed in this paper is the first attempt to achieve high scalability and retain good noise resistance.

3. DENSE SUBGRAPH FINDING REVISIT

In this section, we revisit the dense subgraph finding problem from the perspective of Evolutionary Game Theory (EGT) [35] and discuss the common scalability bottleneck of the infection immunization dynamics (IID) [28] and other affinity-based methods.

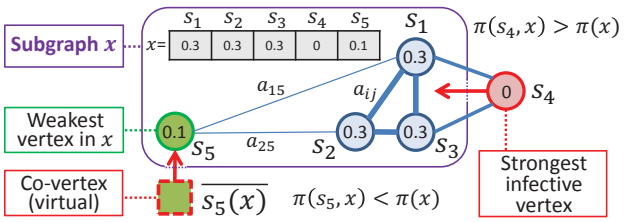

Figure 1: An affinity graph G with 5 vertexes $\{s_1, s_2, s_3, s_4, s_5\}$. Thicker edge means larger affinity. The purple rounded rectangle identifies a subgraph x of G. The subgraph x is composed of vertexes $\{s_1, s_2, s_3, s_5\}$ and is represented by $x = [0.3, 0.3, 0.3, 0, 0.1]^T, x \in \triangle^5$. s_4 is an infective vertex against x, since $\pi(s_4, x) > \pi(x)$. s_5 is a weak vertex in x, since $\pi(s_5, x) < \pi(x)$. The co-vertex $\overline{s_5(x)}$ represents an infective subgraph composed of $\{s_1, s_2, s_3\}$.

Consider a global affinity graph, denoted by $G = (V, I, A)$, where $V = \{v_i \in R^d \mid i \in I = [1, n]\}$ is the set of vertexes, and each vertex v_i uniquely corresponds to a d-dimensional data point in R^d space, R being the set of real numbers, $I = [1, n]$ is the range of indices of all vertexes. A is the affinity matrix, where each entry a_{ij} of A represents the affinity between v_i and v_j, that is,

$$a_{ij} = \begin{cases} e^{-k||v_i - v_j||_p} & i \neq j \\ 0 & i = j \end{cases} \tag{1}$$

where $||\cdot||_p$ represents the L_p-norm ($p \geq 1$) and $k > 0$ is the scaling factor of the *Laplacian Kernel*.

Given an affinity graph G of n vertexes, each graph vertex v_i can be further referenced by an n-dimensional index vector $s_i = [\underbrace{0 \cdots 0}_{i-1} \; 1 \; \underbrace{0 \cdots 0}_{n-i}]^T$. In other words, both v_i and s_i refer to the same i-th vertex in the graph.

A subgraph can be modeled by a subset of vertexes as well as their probabilistic memberships. Take Figure 1 as an example. We assign L_1 normalized non-negative weights to a subgraph x containing vertexes $\{s_1, s_2, s_3, s_5\}$ and represent the subgraph by $x = 0.3 \cdot s_1 + 0.3 \cdot s_2 + 0.3 \cdot s_3 + 0 \cdot s_4 + 0.1 \cdot s_5$. Alternatively, x can be regarded as an n-dimensional vector storing all the vertex weights, where the i-th dimension x_i is the weight of vertex s_i. Intuitively, x_i embodies the probability that vertex s_i belongs to subgraph x. Thus $x_i = 0$ indicates that s_i does not belong to x. For example, vertex s_4 in Figure 1 does not belong to x. In general, a subgraph can be represented by an n-dimensional vector $x \in \triangle^n$ in the standard simplex $\triangle^n = \{x \in R^n \mid \sum_{i \in I} x_i = 1, x_i \geq 0\}$.

The average affinity between two subgraphs $x, y \in \triangle^n$ is measured by the weighted affinity sum between all their member vertexes. That is,

$$\pi(y, x) = y^T A x = \sum_i \sum_j x_i y_j a_{ij} \tag{2}$$

As a special case, when $y = s_i$, that is, y is subgraph of a single vertex, $\pi(s_i, x) = (s_i)^T A x$ represents the average affinity between vertex s_i and subgraph x. Moreover, the average affinity between subgraph x and itself is $\pi(x) = \pi(x, x) = x^T A x$, which measures the internal connection strength between all vertexes of subgraph x. Liu and Yan [19] indicated that such internal connection strength is a robust measurement of the intrinsic cohesiveness of x. Thus, $\pi(x)$ is also called the **density** of subgraph x.

Bulò *et al.* [28] indicated that a dense subgraph is a subgraph of local maximum density $\pi(x)$ [19, 26], and every local maximum argument x^* of $\pi(x)$ uniquely corresponds to a dense subgraph. Thus, the dense subgraph seeking problem can be reduced to the following standard quadratic optimization problem (StQP):

$$\begin{array}{ll} \text{Maximize} & \pi(x) = x^T A x = \sum_i x_i \pi(s_i, x) \\ \text{s.t.} & x \in \triangle^n \end{array} \tag{3}$$

which can be solved by the Infection Immunization Dynamics (IID) [28]. IID finds the dense subgraph x^* by iteratively increasing the graph density $\pi(x)$ in the following two steps.

Infection : for a vertex s_i whose average affinity $\pi(s_i, x)$ is larger than the graph density $\pi(x)$, such as $\pi(s_4, x) > \pi(x)$ in Figure 1, increase the its weight x_i.

Immunization : for a vertex s_i whose average affinity $\pi(s_i, x)$ is smaller than the graph density $\pi(x)$, such as $\pi(s_5, x) < \pi(x)$ in Figure 1, decrease its weight x_i.

For the sake of clarity, we write $\pi(y - x, x) = \pi(y, x) - \pi(x)$ and define the relationship between subgraphs x and y as follow. If $\pi(y - x, x) > 0$, then y is said to be *infective* against x. Otherwise, x is said to be *immune* against y. The set of infective subgraphs against x is defined as:

$$\gamma(x) = \{y \in \triangle^n \mid \pi(y - x, x) > 0\} \tag{4}$$

IID performs **infection** and **immunization** using the following **invasion model**.

$$z = (1 - \varepsilon)x + \varepsilon y \tag{5}$$

The new subgraph $z \in \triangle^n$ is obtained by invading $x \in \triangle^n$ with a subgraph εy, where $\varepsilon \in [0, 1]$ and $y \in \triangle^n$. This transfers an amount of ε weight from the vertexes in subgraph x to the vertexes in subgraph y, while the vertex weights in z still sum to 1. In other words, the weights of vertexes in subgraph x are decreased and the weights of vertexes in subgraph y are increased.

For each iteration, IID selects the optimal graph vertex s_i that maximizes the absolute value of $|\pi(s_i - x, x)|$ by function $s_i = M(x)$, where

$$\begin{array}{l} M(x) = \arg \max_{s_i \in (\mathcal{C}_1 \cup \mathcal{C}_2)} |\pi(s_i - x, x)| \\ \mathcal{C}_1 = \{s_i \mid \pi(s_i - x, x) > 0)\} \\ \mathcal{C}_2 = \{s_i \mid \pi(s_i - x, x) < 0, x_i > 0\} \end{array} \tag{6}$$

If $s_i = M(x) \in \mathcal{C}_1$, then $\pi(s_i, x) > \pi(x)$ and s_i is the strongest infective vertex. In this case, an **infection** is performed by the invasion model (Equation 5) with $y = s_i$. This increases $\pi(x)$ by increasing the weight of infective vertex s_i. For example, in Figure 1, invading x with a weight of ε of the strongest infective vertex s_4 transfers a weight of ε from $\{s_1, s_2, s_3, s_5\}$ to s_4 and increases $\pi(x)$.

If $s_i = M(x) \in \mathcal{C}_2$, then $\pi(s_i, x) < \pi(x)$. Thus x is immune against s_i and s_i is the weakest vertex in subgraph x. In such case, an **immunization** is performed by the invasion model (Equation 5) with $y = \overline{s_i(x)}$:

$$\overline{s_i(x)} = \frac{x_i}{x_i - 1}(s_i - x) + x \tag{7}$$

Here, $\overline{s_i(x)}$ is named the *co-vertex* of s_i, and represents an infective subgraph composed of all the vertexes in subgraph

x except s_i. Thus, invading x with $y = \overline{s_i(x)}$ by the invasion model (Equation 5) reduces the weight of vertex s_i and increases the weights of the other vertexes in subgraph x. For example, in Figure 1, The co-vertex $\overline{s_5(x)}$ represents a subgraph composed of $\{s_1, s_2, s_3\}$. Invading subgraph x with a weight of ε of $\overline{s_5(x)}$ transfers a weight of ε from s_5 to $\{s_1, s_2, s_3\}$ and increases $\pi(x)$.

Formally, the infective vertex (or co-vertex) y of the invasion model (Equation 5) can be selected by $y = S(x)$:

$$S(x) = \begin{cases} s_i & \text{if } s_i = M(x) \in \mathcal{C}_1 \\ \overline{s_i(x)} & \text{if } s_i = M(x) \in \mathcal{C}_2 \end{cases} \tag{8}$$

Bulò *et al.* [28] showed the following.

THEOREM 1 ([28]). *The following three statements are equivalent for $x \in \triangle^n$. 1) x is immune against all vertexes $s_i \in \triangle^n$; 2) $\gamma(x) = \emptyset$; and 3) x is a dense subgraph with local maximum $\pi(x)$.*

According to Theorem 1, IID searches for the global dense subgraph x^* by iteratively shrinking $\gamma(x)$ until $\gamma(x) = \emptyset$.

Moreover, Bulò *et al.* [27] showed the following.

THEOREM 2. *Let $y \in \gamma(x)$ and $z = (1-\varepsilon)x + \varepsilon y$, where $\varepsilon = \varepsilon_y(x)$ is defined as follows.*

$$\varepsilon_y(x) = \begin{cases} \min\left[-\frac{\pi(y-x,x)}{\pi(y-x)}, 1\right] & \text{if } \pi(y-x) < 0 \\ 1 & \text{otherwise} \end{cases} \tag{9}$$

Then, $y \notin \gamma(z)$ and $\pi(z) > \pi(x)$.

According to Theorem 2, any infective subgraph $y \in \gamma(x)$ can be excluded from $\gamma(z)$ by invading x with weight $\varepsilon = \varepsilon_y(x)$ of y. This monotonously reduces the volume of $\gamma(z)$ and guarantees the convergence of IID.

IID needs the full affinity matrix to find the dense subgraph by solving the StQP problem (Equation 3). Therefore, its scalability is largely limited due to the $\mathcal{O}(n^2)$ time and space complexity in computing the complete affinity matrix A. Although forcing A sparse can reduce such expensive time and space cost to some extent [8], the enforced sparsity breaks the intrinsic cohesiveness of dense subgraphs, thus inevitably impairs the detection quality. Both the scalability of DS [26] and SEA [18] are limited due to the same reason, since they need the complete affinity matrix A as well.

4. THE ALID APPROACH

In this section, we introduce our ALID method. The major idea of ALID is to confine the computation of infection and immunization in small local ranges within the Region of Interest (ROI), so that only small submatrices of the affinity matrix need to be computed. As a result, ALID largely avoids the affinity matrix computation and significantly reduces both the time and space complexity.

The framework of ALID is an iterative execution of the following three steps, as summarized in Figure 2.

Step 1 *Finding local dense subgraph by Localized Infection Immunization Dynamics (LID).* We propose a new algorithm LID that confines all infection immunization iterations in a small local range to find the local dense subgraph. LID avoids computing the full affinity matrix by selectively computing a few columns of the local affinity matrix.

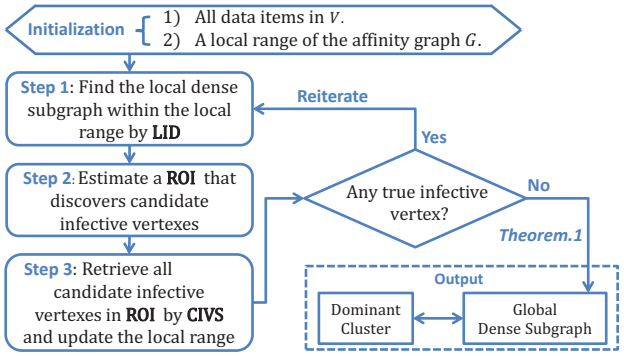
Initialization
1) All data items in V.
2) A local range of the affinity graph G.
Step 1: Find the local dense subgraph within the local range by LID
Reiterate
Yes
Step 2: Estimate a ROI that discovers candidate infective vertexes
Any true infective vertex?
No
Theorem.1
Step 3: Retrieve all candidate infective vertexes in ROI by CIVS and update the local range
Output
Dominant Cluster
Global Dense Subgraph

Figure 2: ALID framework overview.

Step 2 *Estimating a Region of Interest (ROI).* The local dense subgraph found in the LID step may or may not be a global dense subgraph, since there may still be global infective vertexes that are not covered by the current local range, as indicated by Theorem 1. Therefore, we estimate a ROI to identify the candidate infective vertexes in the global range, so that the current local range can be further updated to cover the global dense subgraph.

Step 3 *Candidate Infective Vertex Search (CIVS).* CIVS efficiently retrieves the candidate infective vertexes within a ROI and uses them to update the local range for the next iteration.

The iterative process of ALID terminates when the local dense subgraph found in the last round of iteration is immune against all vertexes in the global range. According to Theorem 1, such a subgraph is a global dense subgraph that identifies a true dominant cluster. We explain the details of the three steps in the first three subsections as follow.

4.1 Localized Infection Immunization Dynamics (Step 1)

The key idea of Localized Infection Immunization Dynamics (LID) is that the dense subgraph on a small local range of the affinity graph can be detected by selectively computing only a few columns of the corresponding local affinity submatrix. Denote by $\beta \subset I$ the local range of the affinity graph, which is the index set of a local group of graph vertexes. LID finds the local dense subgraph $\hat{x}$ that maximizes $\pi(x)$ in the local range β by localizing the infection immunization process on the selectively computed submatrix $A_{\beta\alpha}$ (see Figure 3). Here, $\alpha \triangleq \{i \in \beta \mid x_i > 0\}$ is the *support* of the subgraph $x \in \triangle_\beta^n$, where $\triangle_\beta^n = \{x \in \triangle^n \mid \sum_{i\in\beta} x_i = 1\}$ represents the set of all possible subgraphs within local range β. In this way, the computation of the full matrix $A_{\beta\beta}$ can be effectively avoided. Consequently, LID is more efficient than directly running IID in the local range.

Algorithm 1 shows the steps in a LID iteration, which takes $\left[x^{(t)}, (A_{\beta\alpha}x_\alpha)^{(t)}\right]$ as input and obtains the output $\left[x^{(t+1)}, (A_{\beta\alpha}x_\alpha)^{(t+1)}\right]$ in linear time and space with respect to the size of β. The superscripts (t) and $(t+1)$ indicate the number of iterations. For the interest of simplicity, we omit (t) and illustrate the details of Algorithm 1 as follows.

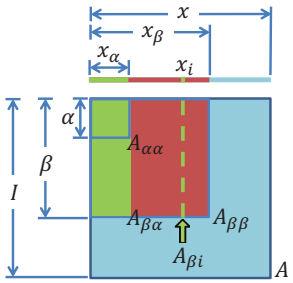

Figure 3: The affinity matrix with ordered data item indexes. A is the global affinity matrix, $A_{\beta\beta}$ is the affinity matrix in the local range β, and $A_{\beta\alpha}$ is the group of columns in $A_{\beta\beta}$ corresponding to the support α. The green dashed line denotes the new column $A_{\beta i}$: $i \in (\beta - \alpha)$. Only the green parts ($A_{\beta\alpha}$ and $A_{\beta i}$) are involved in the LID iteration.

First, we select the infective vertex (or co-vertex) $S(x) \in \triangle_\beta^n$ in the local range β by Equation 6 and 8. Recall that $x \in \triangle_\beta^n$, the component $\pi(s_i - x, x)$ can be computed as

$$\begin{aligned}\pi(s_i - x, x) &= (s_i - x_\beta)^T A_{\beta\beta} x_\beta \\ &= (A_{\beta\alpha}x_\alpha)_i - \sum_{i\in\alpha}(A_{\beta\alpha}x_\alpha)_i x_i\end{aligned} \tag{10}$$

where only the graph vertexes $s_i \in \triangle_\beta^n$ are considered.

Second, we compute the invasion share $\varepsilon_y(x)$ by Equation 9, whose value depends on $y = S(x)$ in two cases:

Case 1 : *Infection* ($y = S(x) = s_i$). The key components $\pi(y - x, x)$ and $\pi(y - x)$ in Equation 9 are computed by Equations 10 and 11, respectively.

$$\begin{aligned}\pi(s_i - x) &= (s_i - x_\beta)^T A_{\beta\beta}(s_i - x_\beta) \\ &= -2(A_{\beta\alpha}x_\alpha)_i + \sum_{i\in\alpha}(A_{\beta\alpha}x_\alpha)_i x_i\end{aligned} \tag{11}$$

Case 2 : *Immunization* ($y = S(x) = \overline{s_i(x)}$). $\varepsilon_y(x)$ is computed by plugging Equation 12 into Equation 9.

$$\begin{cases}\pi(\overline{s_i(x)} - x, x) = \frac{x_i}{x_i - 1}\pi(s_i - x, x) \\ \pi(\overline{s_i(x)} - x) = (\frac{x_i}{x_i - 1})^2 \pi(s_i - x)\end{cases} \tag{12}$$

Last, the new subgraph $x^{(t+1)}$ is obtained by Equation 13 and $(A_{\beta\alpha}x_\alpha)^{(t+1)}$ is computed by Equation 14 in linear time and space for the next iteration.

$$x^{(t+1)} = (1 - \varepsilon_y(x))x + \varepsilon_y(x)S(x) \tag{13}$$

$$(A_{\beta\alpha}x_\alpha)^{(t+1)} = (Ax^{(t+1)})_\beta =$$

$$\begin{cases}\varepsilon_y(x)[A_{\beta i} - A_{\beta\alpha}x_\alpha] + A_{\beta\alpha}x_\alpha & y = s_i \\ (\frac{x_i}{x_i - 1})\varepsilon_y(x)[A_{\beta i} - A_{\beta\alpha}x_\alpha] + A_{\beta\alpha}x_\alpha & y = \overline{s_i(x)}\end{cases} \tag{14}$$

Each LID iteration in Algorithm 1 is guaranteed by Theorem 2 to shrink the size of the local infective subgraph set $\gamma_\beta(x) = \{y \in \triangle_\beta^n \mid \pi(y - x, x) > 0\}$. Thus, we obtain the local dense subgraph $\hat{x} \in \triangle_\beta^n$ in the local range β by repeating Algorithm 1 to shrink $\gamma_\beta(x)$ until $\gamma_\beta(x) = \emptyset$. According to Theorem 1, $\gamma_\beta(\hat{x}) = \emptyset$ indicates that $\hat{x}$ is immune against all vertexes $s_i \in \triangle_\beta^n$, thus $\hat{x}$ is a local dense subgraph with local maximum $\pi(\hat{x})$ in the local range β. In practice, we stop the LID iteration when $\pi(x)$ is stable or the total number of iterations exceeds an upper limit T. Let

Algorithm 1: A single period of LID iteration

Input: $x^{(t)}$, $(A_{\beta\alpha}x_\alpha)^{(t)}$
Output: $x^{(t+1)}$, $(A_{\beta\alpha}x_\alpha)^{(t+1)}$

1: Select the infective vertex $y = S(x)$ by Equation 8
2: Calculate the invasion share $\varepsilon_y(x)$ by Equation 9
3: Update $x^{(t+1)}$ by the invasion model of Equation 13
4: Update $(A_{\beta\alpha}x_\alpha)^{(t+1)}$ by Equation 14
5: **return** $x^{(t+1)}$, $(A_{\beta\alpha}x_\alpha)^{(t+1)}$

$b = |\beta|$ be the size of β, the time and space complexities of LID method are $\mathcal{O}(Tb)$ and $\mathcal{O}(b)$, respectively, not including the time and space cost in computing the affinity matrix A. The initialization of $[x, A_{\beta\alpha}x_\alpha]$ will be discussed in Section 4.3.

All ALID iterations are restricted on the dynamically computed submatrix $A_{\beta\alpha}$, where the new matrix column $A_{\beta i}$ only needs to be computed and stored when $i \in (\beta - \alpha)$ (see Figure 3). We will discuss how LID effectively reduces the original $\mathcal{O}(n^2)$ time and space complexity of affinity matrix computation in Section 4.5.

4.2 Estimating ROI (Step 2)

A local dense subgraph $\hat{x} \in \triangle_\beta^n$ may not be a global dense subgraph in $\triangle^n$, since β may not fully cover the true dense subgraph in the global range I. Therefore, there may still be graph vertexes in the complementary range $U = I - \beta$, which are infective against the local dense subgraph $\hat{x}$. Thus, the current local range β should be updated to include such infective vertexes in U, so that the true dense subgraph with maximum graph density can be detected by LID.

A natural way to find the global dense subgraph x^* is to keep invading $\hat{x}$ using the infective vertexes in U until no infective vertex exists in the global range I. However, fully scanning U for infective vertexes leads to an overall time complexity of $\mathcal{O}(n^2)$ in detecting all dominant clusters, since U contains all the remaining vertexes in I with an overwhelming proportion of irrelevant vertexes. To tackle this problem, we estimate a Region of Interest (ROI) from $\hat{x}$ to include all the infective vertexes and exclude most of the irrelevant ones. Only the limited amount of vertexes inside the ROI are used to update β, which largely reduces the amount of vertexes to be considered and effectively reduces the time and space complexity of ALID.

Before estimating the ROI, we first construct a **double-deck hyperball** $H(D, R_{in}, R_{out})$ from $\hat{x}$, where $D \in R^d$ is the ball center and R_{in}, R_{out} are the radiuses of the inner and outer balls, respectively, which are defined as follows.

$$\begin{cases}D = \sum_{i\in\alpha} v_i\hat{x}_i, \ \textit{where } v_i \in V \textit{ are the data items.} \\ R_{in} = \frac{1}{k}\ln(\frac{\lambda_{in}}{\pi(\hat{x})}), \ \textit{where } \lambda_{in} = \sum_{i\in\alpha}\hat{x}_i e^{-k||v_i - D||_p} \\ R_{out} = \frac{1}{k}\ln(\frac{\lambda_{out}}{\pi(\hat{x})}), \ \textit{where } \lambda_{out} = \sum_{i\in\alpha}\hat{x}_i e^{k||v_i - D||_p}\end{cases} \tag{15}$$

where k is the positive scaling factor in Equation 1. Proposition 1 gives two important properties of the double-deck hyperball $H(D, R_{in}, R_{out})$.

According to the two properties in Proposition 1, the surfaces of the inner ball $H(D, R_{in})$ and the outer ball $H(D, R_{out})$ are two boundaries, which guarantee that every

data entry inside the inner ball corresponds to an infective vertex and the ones outside the outer ball are non-infective.

PROPOSITION 1. *Given the local dense subgraph $\hat{x}$, the double-deck hyperball $H(D, R_{in}, R_{out})$ has the following properties:*

1. $\forall j \in I$ *and* $||v_j - D||_p < R_{in}$, $\pi(s_j - \hat{x}, \hat{x}) > 0$; *and*
2. $\forall j \in I$ *and* $||v_j - D||_p > R_{out}$, $\pi(s_j - \hat{x}, \hat{x}) < 0$.

Proposition 1 is proved in Section A of appendix in [10].

The **ROI** is defined as a growing hyperball $H_c(D, R)$, whose surface starts from the inner ball and gradually approaches the outer ball as the ALID iteration continues. The radius R is defined as

$$R = R_{in} + \theta(c)(R_{out} - R_{in}) \tag{16}$$

where $\theta(c) = \frac{1}{1+e^{(4-c/2)}}$ is a shifted logistic function to control the growing speed and c is the current number of ALID iteration. When c grows large, we have $\theta(c) \approx 1$ and $R \approx R_{out}$, thus, the ROI is guaranteed to coincide with the outer ball as c grows. Since the outer ball $H(D, R_{out})$ is guaranteed by Proposition 1 to contain all infective vertexes in the global range I, the finally found local dense subgraph $\hat{x}$ within the ROI is guaranteed to be immune against all vertexes in I, thus $\hat{x}$ is a global dense subgraph according to Theorem 1. Moreover, starting the ROI from the small inner ball can effectively reduce the number of vertexes to be scanned in the first several ALID iterations.

4.3 Candidate Infective Vertex Search (Step 3)

The hyperball of ROI identifies a small local region in the d-dimensional space R^d. The data items $v_i \in V$ inside the ROI correspond to the candidate graph vertexes, which are probably infective against the current local dense subgraph $\hat{x}$ and may further increase the graph density $\pi(\hat{x})$. Therefore, we carefully design the Candidate Infective Vertex Searching (CIVS) method to efficiently retrieve such data items inside the ROI and use them to update the local range $\beta^{(c)}$. The variable c is the current number of ALID iteration.

Retrieving the data items inside the ROI $H_c(D, R)$ is equivalent to a *fixed-radius near neighbor problem* in the d-dimensional space R^d. For the interest of scalability, we adopt the Locality Sensitive Hashing (LSH) method [12] to solve this problem. However, LSH can only retrieve data items within a small Locality Sensitive Region (LSR) of a query (Figure 4). Thus, when using the hyperball center D as a single LSH query (Figure 4 (a)), the corresponding LSR may fail to find all data items within the ROI, which would prevent ALID from converging to the optimal result.

To solve this problem, we propose CIVS to guarantee that ALID converges to the optimal result. The convergence of ALID is proved in Section B of the appendix in [10]. CIVS applies multiple LSH queries using all the supporting data items of $\hat{x}^{(c)}$ (i.e., $\{v_i \in V \mid \hat{x}_i^{(c)} > 0\}$). As shown in Figure 4 (b), the benefit of CIVS is that multiple LSRs effectively cover most of the ROI, thus most of the data items within the ROI can be retrieved. Specifically, CIVS first collect all the new data items that is retrieved by any of the supporting data items of $\hat{x}^{(c)}$, using LSH. Then, we retrieve at most δ new data items within the ROI that are the nearest to the ball center D. The index set of the retrieved data items is denoted by $\psi = \{i \mid i \in (I - \alpha^{(c)}), v_i \in H_c(D, R)\}, |\psi| \leq \delta$.

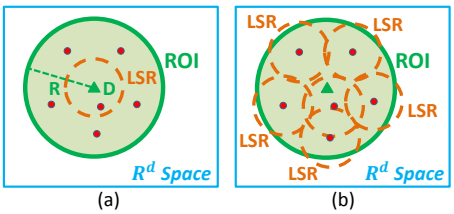

Figure 4: Each dashed circle is a locality sensitive region (LSR). The red points show the supporting data items of the current local dense subgraph $\hat{x}^{(c)}$. (a) A single LSR cannot cover the ROI. (b) CIVS covers most of the ROI by multiple LSRs.

The retrieved data $\{v_i \mid i \in \psi\}$ are used to perform update: $\left[\hat{x}^{(c)}, (A_{\beta\alpha}\hat{x}_\alpha)^{(c)}\right] \rightarrow \left[x^{(c+1)}, (A_{\beta\alpha}x_\alpha)^{(c+1)}\right]$ as follow.

$$\begin{cases} x^{(c+1)} = \hat{x}^{(c)} \\ (A_{\beta\alpha}x_\alpha)^{(c+1)} = \begin{bmatrix} (A_{\alpha\alpha}\hat{x}_\alpha)^{(c)} \\ (A_{\psi\alpha}\hat{x}_\alpha)^{(c)} \end{bmatrix} \end{cases} \tag{17}$$

where $(A_{\psi\alpha}\hat{x}_\alpha)^{(c)} = A_{\psi\alpha}\hat{x}_\alpha^{(c)}$

Equation 17 updates the local range by $\beta^{(c+1)}{=}\alpha^{(c)} \cup \psi$, where ψ involves new infective vertexes against $x^{(c+1)}$. Then, we can re-run LID (i.e., Step 1) with the initialization of $[x^{(c+1)}, (A_{\beta\alpha}x_\alpha)^{(c+1)}]$ to find the local dense subgraph $\hat{x}^{(c+1)}$ in the new range $\beta^{(c+1)}$. Since $\hat{x}^{(c+1)}$ is guaranteed by Theorem 1 to be immune against all vertexes in $\psi \subset U$, the number of infective vertexes in the global range I is further reduced.

The time and space complexity for building the hash tables are linear with respect to n. Specifically, the time complexity to build l hash tables by μ hash functions is $\mathcal{O}(ndl\mu)$. The space complexity consists of $\mathcal{O}(nd)$ space for all the d dimensional data items, $\mathcal{O}(nl)$ space for an inverted list that maps each data item to their buckets and $\mathcal{O}(nl)$ space for l hash tables [12]. Since all possible LSH queries are built into the hash tables, we check the inverted list to retrieve neighbor data items and do not store the hash keys.

4.4 Summarization of ALID

The entire iteration of ALID is summarized in Algorithm 2. The LID in Step 1 makes the local dense subgraph immune against all vertexes within a local range of the ROI. The ROI and CIVS in Step 2 and Step 3 update the local range by the new infective vertexes retrieved from global range. In this way, the number of infective vertexes in global range is guaranteed to be iteratively reduced to zero. Then, according to Theorem 1, the last found local dense subgraph is a global one that identifies a dominant cluster [35]. Algorithm 2 stops when a global dense subgraph is found or the total number of iterations exceeds an upper limit C. Since Algorithm 2 is initialized with $A_{\beta\alpha}x_\alpha = 0$, which cannot be used to compute the radius of ROI $H_{c=1}(D, R)$ (Equation 15 and 16), thus we set $R = 0.4$ for the first iteration $c = 1$.

In order to fairly compare with the other affinity-based methods, ALID adopts the same *peeling method* as DS [26] and IID [28] do to detect all dominant clusters. The *peeling method* peels off the detected cluster and reiterates on the remaining data items to find another one until all data items are peeled off. Then, the clusters with large values of $\pi(x)$ (e.g., $\pi(x) \geq 0.75$) are selected as the final result.

Algorithm 2: The entire ALID iteration

Input: An initial vertex index $i \in I$
Output: A global dense subgraph x^* in global range I

1: **Set** $\alpha = \beta = i$, $x = s_i$, $A_{\beta\alpha}x_\alpha = a_{ii} = 0$, $c = 1$
2: **repeat**
3: **Step 1**: $\left[x^{(c)}, (A_{\beta\alpha}x_\alpha)^{(c)}\right] \rightarrow \left[\hat{x}^{(c)}, (A_{\beta\alpha}\hat{x}_\alpha)^{(c)}\right]$
Find the local dense subgraph $\hat{x}^{(c)}$ by LID in Step 1
4: **Step 2**: $\hat{x}^{(c)} \rightarrow H_c(D, R)$
Estimate ROI from the local dense subgraph $\hat{x}^{(c)}$
5: **Step 3**: $\left[\hat{x}^{(c)}, (A_{\beta\alpha}\hat{x}_\alpha)^{(c)}\right] \rightarrow \left[x^{(c+1)}, (A_{\beta\alpha}x_\alpha)^{(c+1)}\right]$
Apply CIVS to retrieve candidate vertexes within the ROI and update the local range $\beta^{(c)}$ by Equation 17 for the next iteration
6: Index update: $c \leftarrow c + 1$
7: **until** $\hat{x}^{(c)}$ is a global dense subgraph, or $c > C$
8: **return** $x^* = \hat{x}^{(c)}$

4.5 Complexity Analysis

The time and space complexities of ALID mainly consist of three parts:

1) The time and space complexities of LID in Step 1 are $\mathcal{O}(Tb)$ and $\mathcal{O}(b)$, respectively, where $b = |\beta| < n$ is the size of the local range β and T is the constant limit of the number of LID iterations.
2) The time and space complexities for the hash tables of CIVS are $\mathcal{O}(ndl\mu)$ and $\mathcal{O}(n(2l + d))$, respectively, where d, l, μ are constant LSH parameters.
3) The time and space complexities for the affinity matrix A are $\mathcal{O}(C(a^* + \delta)n)$ and $\mathcal{O}(a^*(a^* + \delta))$, respectively, which are analyzed in detail as follows.

Since all ALID iterations are restricted by $A_{\beta\alpha}$, the time and space complexities for the affinity matrix are determined by the size of $A_{\beta\alpha}$. Let $(A_{\beta\alpha})_i^c$ be the submatrix computed in the c-th iteration in Algorithm 2 when detecting the i-th cluster. Denote by a_i^c and b_i^c, respectively, the column and row sizes of $(A_{\beta\alpha})_i^c$. Since the maximum number of iterations of Algorithm 2 is C, the overall time cost for detecting the i-th cluster is $Time(i) < \sum_{c=1}^{C} a_i^c b_i^c$, which is a *loose bound* since many matrix entries of different $(A_{\beta\alpha})_i^c$ are duplicate and only computed once. Then, we can derive

$$Time(i) < Ca_i^C(a_i^C + \delta) \tag{18}$$

from the following observations:

- $a_i^c \leq a_i^C$, since more and more matrix columns (i.e., $A_{\beta i} : i \in (\beta - \alpha)$ in Figure 3) are computed.
- $b_i^c \leq (a_i^c + \delta) \leq (a_i^C + \delta)$, since the size of β is strictly limited by the ROI, where at most δ data items can be retrieved by CIVS.

Since $Time(i)$ is the time cost of affinity matrix computation for detecting the i-th dominant cluster, then the overall cost in time of detecting all dominant clusters is $\sum_i Time(i)$. We can derive from Equation 18 that

$$\sum_i Time(i) < \sum_i Ca_i^C(a_i^C + \delta) \tag{19}$$

Recall that ALID adopts the *peeling method* (see Section 4.4), which peels off one detected cluster and reiterates on the remaining data items to find another one until all the n data items are peeled off. Therefore, we can derive

$$\sum_i a_i^C = n \tag{20}$$

from the fact that a_i^C is the size of the i-th detected cluster.

Define $a^* = \max_i\{a_i^C\}$ and $b^* = \max_i\{b_i^C\}$, where $b^* \leq (a^* + \delta)$ due to the restriction of ROI. Then, we can derive from Equation 19 and Equation 20 that the overall cost in time of computing the affinity matrix is

$$\sum_i Time(i) < \sum_i Ca_i^C(a^* + \delta) = C(a^* + \delta)n \tag{21}$$

The maximum cost in space is $a^*b^* \leq a^*(a^* + \delta)$, since all submatrices $(A_{\beta\alpha})_i^c$ are released when the i-th cluster is peeled off. As a result, the time and space complexities for the affinity matrix of ALID are $\mathcal{O}(C(a^*+\delta)n)$ and $\mathcal{O}(a^*(a^*+\delta))$, respectively.

Recall that $a^* = \max_i\{a_i^C\}$ is the size of the largest (single) dominant cluster. We summarize in Table 1 the three typical cases how a^* affects the time and space complexities of the affinity matrix. The data items belonging to the largest dominant cluster with size a^* is referred to as "positive data", data items that do not belong to any dominant cluster are regarded as "noise data" and the size of the entire data set is denoted by n.

Table 1: The complexity of the affinity matrix

Typical Cases	Time Complexity	Space Complexity
$a^* = \omega n$ $(\omega \leq 1)$	$\mathcal{O}(C(\omega n^2 + \delta n))$	$\mathcal{O}(\omega^2 n^2 + \delta\omega n)$
$a^* = n^\eta$ $(\eta < 1)$	$\mathcal{O}(C(n^{1+\eta} + \delta n))$	$\mathcal{O}(n^{2\eta} + \delta n^\eta)$
$a^* \leq P$	$\mathcal{O}(C((P + \delta)n))$	$\mathcal{O}(P^2 + \delta P)$

First, for clean data source, the amount of positive data is in constant proportion of the entire data set. Thus, we have $a^* = \omega n$ and $\omega \leq 1$ is the constant proportion. In this case, ALID reduces the original $\mathcal{O}(n^2)$ time and space complexities of the affinity matrix to $\mathcal{O}(C(\omega n^2 + \delta n))$ and $\mathcal{O}(\omega^2 n^2 + \delta\omega n)$, respectively.

Second, for noisy data source (e.g., tweet-streams and user comments) that generates noise data faster than positive data, the growth rate of a^* is slower than n. In this case, we have $a^* = n^\eta$ $(\eta < 1)$, thus ALID reduces the $\mathcal{O}(n^2)$ time and space complexities to $\mathcal{O}(C(n^{1+\eta} + \delta n))$ and $\mathcal{O}(n^{2\eta} + \delta n^\eta)$, respectively.

Third, for noisy data source with size-limited dominant clusters, there is a constant upper bound P for a^* (i.e., $a^* \leq P$). Typical data sources of such kind consist of phone books and email contacts, where the largest number of people in a stable social group (i.e., dominant cluster) is limited by the *Dunbar's number* [1] [14]. In such a case, we have $a^* \leq P$, thus, the time and space complexities of the affinity matrix are reduced to $\mathcal{O}(C(P+\delta)n)$ and $\mathcal{O}(P^2+\delta P)$, respectively. Note that, since the infection immunization process converges quickly in finite steps [27], a small value of $C = 10$ is adequate for the convergence of ALID.

[1] *Dunbar's number* was proposed by Robin Dunbar, who found that the social group size is upper bounded due to the limited brain size of human. Therefore, the size of a single dominate cluster (i.e., social group) in a social data source is upper bounded by a constant. `http://en.wikipedia.org/wiki/Dunbar's_number`

We conducted extensive experiments on three synthetic data sets to simulate the three typical cases of $a^* = \omega n$, $a^* = n^\eta$ and $a^* \leq P$. The experimental results in Figure 7 of Section 5.2 are consistent with the time and space complexities summarized in Table 1.

4.6 Parallel ALID

ALID is very suitable for parallelization in the MapReduce framework [13], since multiple tasks of ALID can be concurrently run in independent local ranges of the affinity graph. We introduce the parallel ALID (PALID) in Algorithm 3 and provide an illustrative example in Figure 5.

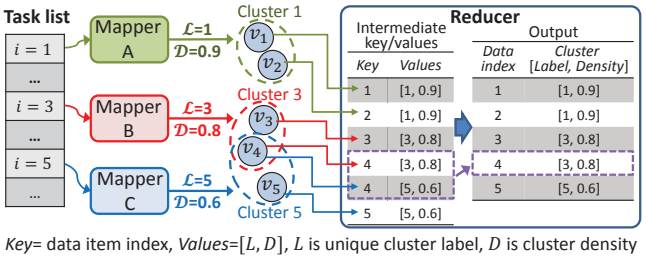

Figure 5: An illustrative example of Algorithm 3. Each mapper runs Algorithm 2 independently with a different initial vertex. The reducer assigns each data item to the cluster with maximum density.

As shown in Figure 5, three initial graph vertex indexes $i = \{1, 3, 5\}$ are assigned to three Mappers (A,B,C). Each Mapper runs Algorithm 2 to detect a cluster independently. Once a Mapper detects a cluster, it produces a list of intermediate key/value pairs ($Key, Values$=$[\mathcal{L}, \mathcal{D}]$). Key is the index of one data item belonging to the cluster, $\mathcal{L}$ is the unique cluster label for the detected cluster and $\mathcal{D}$ is the density of the cluster. In case of overlapping clusters, such as clusters 3 and 5 in the figure, we simply assign the overlapped data item v_4 to cluster 3 of maximum density, which can be easily handled by a reducer.

Since data items belonging to the same dominant cluster are highly similar with each other, such data items are likely to be mapped to the same set of LSH buckets. Therefore, large-sized buckets reveal the potential data items of dominant clusters. As a result, PALID uniformly samples the initial graph vertexes from every LSH hash bucket that contains more than 5 data items. The sample rate is 20%.

The hash tables and data items are stored in a server database and accessed via the network. The communication overhead is low due to the following reasons: 1) the LSH queries only transport data item indexes, which consumes ignorable bandwidth; 2) each mapper only needs to access a few data items to detect one dominant cluster due to the local property of ALID.

In summary, ALID is highly parallelizeable in the MapReduce framework, which further improves its capability in handling massive data in real world applications.

5. EXPERIMENTAL RESULTS

In this section, we empirically examine and analyze the performances of ALID and PALID. The following state-of-the-art affinity-based methods are analyzed as well: 1) Affinity Propagation (AP) [16]; 2) the Shrinking Expansion Algorithm (SEA) [18]; 3) Infection Immunization Dynamics

Algorithm 3: The parallel ALID (PALID)

Input: $V = \{v_i \mid i \in I = [1, n]\}$
Output: Cluster labels and cluster densities of each data item in the detected clusters
Tasklist: A list of initial graph vertex indexes
Map(Key, $Value$)
\\Key: An initial vertex index i for Algorithm 2
\\$Value$: A unique cluster label $\mathcal{L}$ for the detected cluster
1: Call Algorithm 2 to find a global dense subgraph x^*, which identifies dominant cluster $\mathcal{L}$ with density $\mathcal{D} = \pi(x^*)$
2: **for all** data item index h in $I = [1, n]$ **do**
3: **if** $x_h^* > 0$ **then**
4: Emit(h, $[\mathcal{L}, \mathcal{D}]$) \\ v_h belongs to cluster $\mathcal{L}$
5: **end if**
6: **end for**
Reduce(Key, $Values$)
\\Key: The index h of a single data item $v_h \in V$
\\$Values$: A list of $[\mathcal{L}, \mathcal{D}]$ w.r.t. the clusters containing v_h
1: Find $[\mathcal{L}^*, \mathcal{D}^*]$ in $Values$ with maximum density $\mathcal{D}^*$
2: Emit(h, $[\mathcal{L}^*, \mathcal{D}^*]$) \\ Assign v_h to cluster $\mathcal{L}^*$

(IID) [28]. We use the published source codes of AP [1] and SEA [2]; as the code for IID is unavailable, we implemented it in MATLAB. All compared methods are carefully tuned to their best performances. The parameter δ in Step 3 of ALID (PALID) is fixed as $\delta = 800$.

The detection quality is evaluated by the Average F_1 score (AVG-F), which is the same criterion as Chen *et al.* [7] used. AVG-F is obtained by averaging the F_1 scores on all the true dominant clusters. A higher F_1 score means a smaller deviation between the detected and the true dominant clusters, thus indicates a better detection quality. Besides, as Chen *et al.* [7] showed, since the data items are partially clustered in this task, traditional evaluation criteria, such as entropy and normalized mutual information, are not appropriate in evaluating the detection quality.

The detection efficiency is measured from two perspectives: 1) the runtime of each method including the time to compute the affinity matrix; 2) the memory overhead, including the memory to store the affinity matrix.

We use a PC computer with a Core i-5 CPU, 12 GB main memory and a 7200 RPM hard drive, running Microsoft Windows 7 operating system. All experiments using single computer are conducted using *MATLAB*, since the methods to be compared were implemented on the same platform. We will explicitly illustrate the parallel experiments of PALID in Section 5.3. PALID was implemented in Java on *Apache Spark* to efficiently process up to 50 million data items.

The following data sets are used: 1) the news articles data set (NART); 2) the near duplicate image data set (NDI); 3) three synthetic data sets; 4) the SIFT-50M data set that consists of 50 million SIFT features [22]. Details about the three synthetic data sets and the SIFT-50M data set are illustrated in Section 5.2 and Section 5.3, respectively. Details of NART and NDI are illustrated as follow. For all data sets, the pairwise distance and affinity are calculated using Euclidean distance and Equation 1 ($p = 2$), respectively.

The news articles data set (NART) is built by crawling 5,301 news articles from `news.sina.com.cn`. It contains 13 real world "hot" events happened from May 2012 to June

2012, each of which corresponds to a dominant cluster of news articles. All 734 news articles of the 13 dominant clusters are manually labeled as ground truth by 3 volunteers without professional background. The remaining 4,567 articles are daily news that do not form any dominant cluster. Each article is represented by a normalized 350-dimensional vector generated by standard Latent Dirichlet Allocation (LDA) [6].

The near duplicate image data set (NDI) contains 109,815 images crawled from `images.google.com.hk`. It includes a labeled set of ground truth of 57 dominant clusters and 11,951 near duplicate images, where images with similar contents are grouped as one dominant cluster. The remaining 97,864 images with diverse contents are regarded as background noise data. Each image is represented by a 256-dimensional GIST feature [25] that describes the global texture of the image content.

5.1 Sparsity Influence Analysis

As mentioned in Section 2, the scalability of canonical affinity-based methods (i.e., IID, SEA, AP) is limited by the $\mathcal{O}(n^2)$ time and space complexity to fully compute and store the affinity matrix. Although the computational efficiency can be improved by sparsifying the affinity matrix [8], the enforced sparsity also breaks the high cohesiveness of dense subgraphs, which inevitably weakens the noise resistance capability and impairs the detection quality.

In this section, we specifically analyze how the sparse degree of sparsified affinity matrix affects the detection quality and runtime of all compared methods. The sparse degree is defined as the ratio of the number of entries in the matrix taking value 0 over the total number of entries in the matrix.

All experiment results are obtained on NART and Sub-NDI data sets. Sub-NDI is a subset of the NDI data set, it contains 6 clusters of 1420 ground truth images and 8520 noise images. We use Sub-NDI instead of NDI, since AP cannot deal with the entire NDI data set with 12GB RAM.

Chen *et al.* [8] provided two approaches to sparsify the affinity matrix: the exact nearest neighbors (ENN) method and the approximate nearest neighbors (ANN) method. The ENN method is expensive on large data sets, while the ANN method can be efficient by employing LSH [12] and Spill-Tree [20]. In our experiments, we sparsify the affinity matrix by LSH due to its efficiency.

For AP, IID and SEA, we directly apply LSH to sparsify the affinity matrix, where only the affinities between the nearest neighbors are computed and stored. The same LSH module is utilized by CIVS in ALID. To remove possible uncertainties caused by the LSH approximation, the parameter settings of LSH are kept exactly the same for all the compared methods, including ALID.

Standard LSH projects each data item onto an equally segmented real line. The line segment length r controls the recall of LSH, and thus affects the sparse degree of the affinity matrix. Figure 6 shows that the sparse degree of all affinity-based methods, including ALID, decreases when r increases. However, the sparse degree of ALID remains high, since ALID only computes small submatrices corresponding to the vertexes within the ROI.

Figures 6(a) and 6(c) show the experiment results on data set NART. The AVG-F of all methods increases to a stable level as sparse degree decreases. This is because the cohesiveness of dense subgraphs are better retained as

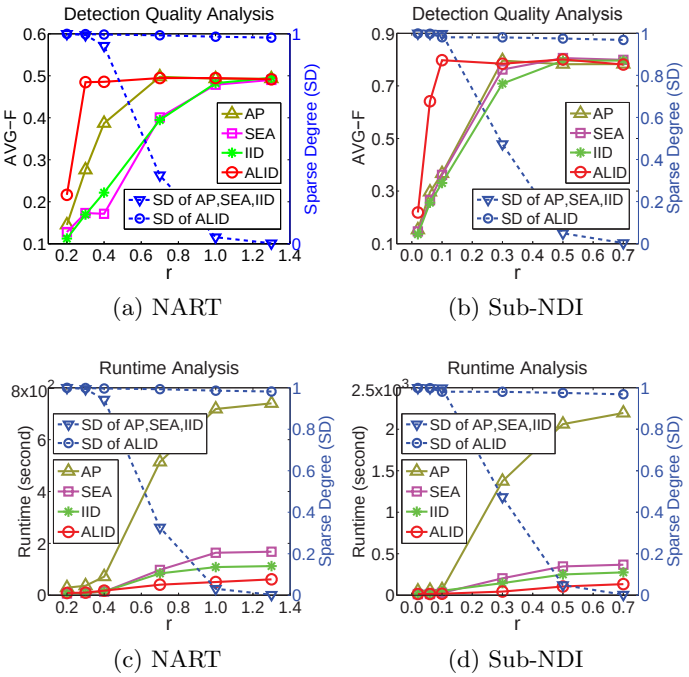

(a) NART (b) Sub-NDI

(c) NART (d) Sub-NDI

Figure 6: The results on NART and Sub-NDI. (a)-(b) How sparse degree affects AVG-F. (c)-(d) How sparse degree affects runtime. For LSH, we use 40 projections per hash value and 50 hash tables. r is the length of the equally divided segments of LSH.

sparse degree decreases. For AP, SEA and IID, when sparse degree approaches zero, the original subgraph cohesiveness are maximally preserved by a full affinity matrix, and thus they all approach their best performances. Since most dense subgraphs exist in small local ranges, the relative local submatrices are good enough to retain their cohesiveness. ALID largely preserves such cohesiveness by accurately estimating the local range of true dense subgraphs and fully computing the relative local affinity matrices. Consequently, ALID achieves a good AVG-F performance under an extremely high sparse degree of 0.998 ($r = 0.3$), which indicates that the calculation and storage of 99.8% matrix entries are effectively pruned. Such situation is rationale, since the useful matrix entries that correspond to the 13 true clusters of 734 data items in data set NART only take $734^2/(13 \times 5301^2) = 0.147\%$ of the entire affinity matrix.

The results in Figure 6(c) demonstrates that sparsifying the affinity matrix reduces the runtime of the affinity-based methods. The runtime of all methods are comparably low when sparse degree is high. However, when sparse degree decreases, the differences in runtime among the methods become significant. When $r = 1.3$, AP is significantly slower than the other methods due to its expensive message passing overheads. Moreover, SEA is much slower than IID due to the time consuming replicator dynamics [35]. ALID is the fastest, since it effectively prunes the computation of 99.8% affinity matrix entries. Similar results are also observed on data set Sub-NDI (Figures 6(b) and 6(d)).

5.2 Scalability Analysis

In this section, we analyze the scalability of the affinity-based methods on four data sets, which consist of the real

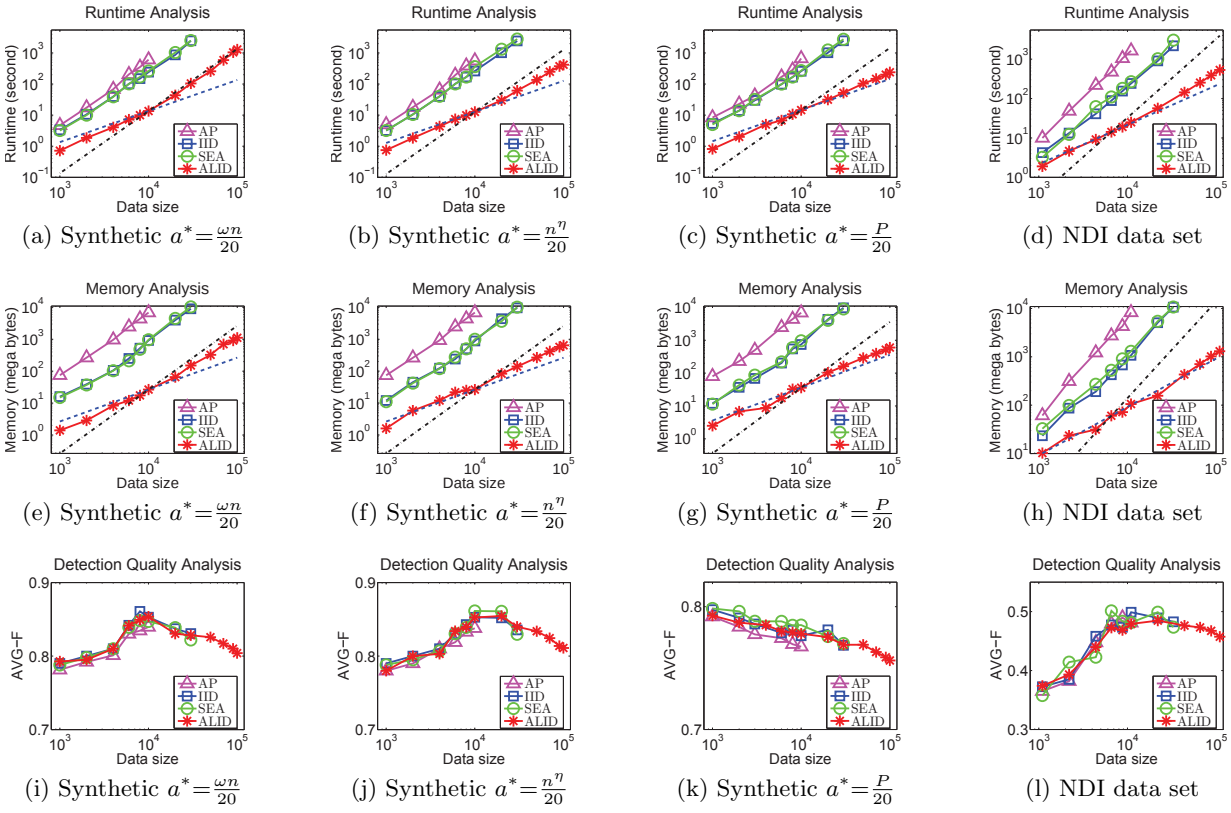

(a) Synthetic $a^*=\frac{\omega n}{20}$ (b) Synthetic $a^*=\frac{n^\eta}{20}$ (c) Synthetic $a^*=\frac{P}{20}$ (d) NDI data set

(e) Synthetic $a^*=\frac{\omega n}{20}$ (f) Synthetic $a^*=\frac{n^\eta}{20}$ (g) Synthetic $a^*=\frac{P}{20}$ (h) NDI data set

(i) Synthetic $a^*=\frac{\omega n}{20}$ (j) Synthetic $a^*=\frac{n^\eta}{20}$ (k) Synthetic $a^*=\frac{P}{20}$ (l) NDI data set

Figure 7: Scalability analysis on synthetic data sets with 3 typical cases of a^* and the real world data set NDI. Parameters for the 3 synthetic data sets are: $\omega = 1.0$, $\eta = 0.9$ and $P = 1000$. The blue and black dashed lines are reference lines with slope 1 and 2, respectively. Under double logarithmic coordinate system, the slope of performance curves indicate the orders of growth with respect to the size of data set.

world data set NDI and the other three synthetic data sets.

The *synthetic data sets* are made up by sampling n 100-dimensional data items from 20 different multivariate gaussian distributions as dominant clusters and one uniform distribution as the background noise. To better simulate typical properties of real world data, we make some gaussian distributions partially overlapped by setting their mean vectors close to each other and variate the shapes of all gaussian distributions by different diagonal covariance matrices with elements ranged in $[0, 10]$. Then, we sample a^* data items from each gaussian distribution as ground truth and $(n - 20a^*)$ data items from the surrounding uniform distribution as noise. Since all the 20 clusters are sampled in equal size, a^* is the largest size of dominant clusters, which is consistent with the definition of a^* in Section 4.5.

The *synthetic data sets* are mainly used to test the efficiency and scalability of ALID, thus we simulate the three typical cases of a^* analyzed in Table 1 by controlling the amount of sampled data items with $a^*=\frac{\omega n}{20}$, $a^*=\frac{n^\eta}{20}$ and $a^*=\frac{P}{20}$, where the constant denominator 20 does not affect the complexity of ALID. We use different values of $n \in [1\times10^3, 1\times10^5]$ to generate data sets of different sizes. For the experiment on the NDI data set, we generate subsets of different sizes by randomly sampling the original NDI data set. The performances are evaluated in AVG-F, runtime and memory overheads of methods SEA, IID, AP and ALID.

We adopt the double logarithmic coordinate system to draw all the performance curves of runtime and memory overheads. In this way, the empirical orders of runtime growth with respect to the data set size (i.e., n) can be easily observed from the slopes of the performance curves. Take $runtime = n^2$ as an example, the slope of the runtime curve under double logarithmic coordinate system is $\log(runtime)/\log(n) = 2$, which is consistent with the theoretical order of runtime growth. We can observe the empirical order of memory growth in the same way.

As it is shown in Figures 7(a)-7(c), the empirical orders of runtime growth on the synthetic data sets are $\log(runtime)/\log(n) \approx 2$ when $a^*=\frac{\omega n}{20}$ ($\omega = 1.0$), $\log(runtime)/\log(n) \approx 1.7$, when $a^*=\frac{n^\eta}{20}$ ($\eta = 0.9$) and $\log(runtime)/\log(n) \approx 1$, when $a^*=\frac{P}{20}$, respectively. This result is consistent with the orders of time complexity (see Table 1) analyzed in Section 4.5. The results on the real world data set NDI (Figure 7(d)) also demonstrate that the empirical order of runtime growth of ALID is is substantially lower than the other affinity-based methods.

As shown in Figure 7(e)-7(g), the empirical order of memory growth of ALID on the three typical synthetic data sets are consistent with the theoretical space complexity (see Table 1) analyzed in Section 4.5. The results on NDI (Figure 7(h)) further demonstrates the superior memory performance of ALID, which consumes about 90% less memory to process 3.3 times larger data size with a substantially lower empirical order of memory growth than the other

affinity-based methods.

The experimental results in Figures 7(i)-7(l) show that ALID achieves comparable AVG-F performance with the other affinity-based methods on all the four data sets. The AVG-F curves in Figures 7(i),7(j),7(l) have similar trends as the data set size increases: the AVG-F grows first due to the increasing dense subgraph cohesiveness caused by the growing amount of ground truth data (i.e., $a^*=\frac{\omega n}{20}$ and $a^*=\frac{n^\eta}{20}$). Then, the AVG-F decreases due to the growing amount of ground truth data in overlapping dominant clusters and the increasing influence of noise. The AVG-F curves in Figure 7(k) monotonously decreases, since the dense subgraph cohesiveness does not increase when the amount of ground truth data is fixed to $a^*=\frac{P}{20}$.

In summary, ALID achieves remarkably better scalability than the other affinity-based methods, and at the same time retains high detection quality. The high scalability is mainly achieved by limiting all ALID iterations within the ROI, which largely prevents unnecessary storage and computation of the entire affinity matrix.

5.3 Parallel Experiments of PALID

PALID is the parallel implementation of ALID. It was implemented in Java on the parallel platform of *Apache Spark* (`http://spark.apache.org/`) on operating system Ubuntu. In this section, we evaluate the parallel performance of PALID on a cluster of 5 PC computers each with an i7 CPU, 32 GB RAM and a 7200 RPM hard drive. We set 1 machine as the master and the other 4 machines as workers. Each worker is assigned with 2 executor processes, where each executor takes a single CPU core. The hash tables and the data items are stored in a MongoDB (`http://www.mongodb.org/`) server on the master. All machines are connected by a 1000 Mbit/s ethernet. The performance of PALID is evaluated on the SIFT-50M data set, which is an unlabeled data set containing 50 million SIFT features [22]. The SIFTs are extracted from the IPDID/1000k image data set [9] using the VLFeat toolbox [32].

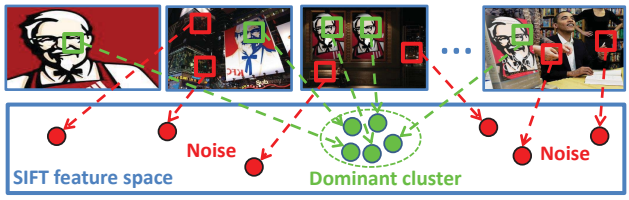

Figure 8: Illustration of SIFT dominant cluster.

Scale Invariant Feature Transform (SIFT) [22] is a L_2 normalized 128-dimensional vector describing the texture of a small image region. In the field of computer vision, it is a standard procedure to represent similar image regions by highly cohesive SIFT dominant clusters named "*visual words*" [31]. As shown in Figure 8, since partial duplicate images always share a common image content (e.g., KFC grandpa), the SIFTs [22] extracted from similar image regions are highly similar to each other and naturally form a dominant cluster (i.e., visual word). However, the number of visual words is unknown and there are also a large proportion of noisy SIFTs extracted from the random non-duplicate regions (i.e., red points in Figure 8 and Figure 10), leading to a high degree of background noise. As a result, the scalability and strong resistance against noise of PALID are very suitable for visual word generation.

Table 2: Performance of PALID on SIFT-50M

Methods	Executors	Runtime	Speadup Ratio
PALID-1Exec	1	17.2 hours	1
PALID-2Exec	2	8.96 hours	1.92
PALID-4Exec	4	4.48 hours	3.84
PALID-8Exec	8	2.29 hours	7.51

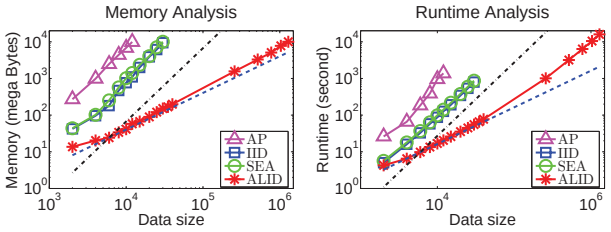

Figure 9: Scalability analysis on SIFT-50M subset.

Table 2 shows the parallel performance of PALID. PALID is able to process 50 million SIFT features in 2.29 hours, achieving a speedup ratio of 7.51 with 8 executors. This demonstrates the promising parallel performance of PALID.

Since there is no parallel solution for the other affinity-based methods, such as IID[28], SEA[18] and AP[16], we fairly compare them with ALID on the SIFT-50M data set, using the same single-machine experimental settings as Section 5.2 used. Figure 9 shows the memory and runtime performances of the affinity-based methods on uniformly sampled subsets of SIFT-50M, where all experiments are stopped when the 12GB RAM limit is reached. We can see that the empirical orders of runtime and memory growth of ALID are significantly lower than the other methods. Especially, ALID consumes 10 GB memory to process 1.29 million SIFTs in 4.4 hours on a standard PC. In contrast, such a large amount of data is far beyond the capability of the other affinity-based methods, which can at most deal with 0.04 million SIFTs on the same platform.

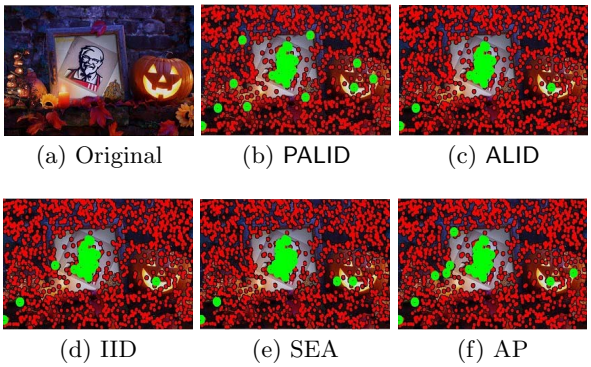

(a) Original (b) PALID (c) ALID

(d) IID (e) SEA (f) AP

Figure 10: Detection quality on (a) "KFC grandpa" of the IPDID/1000k data set. The green points are SIFTs from dominant clusters with high densities ($\pi(x) > 0.75$). The red points in (b)-(f) are noise SIFTs filtered out by the affinity-based methods.

Despite that SIFT-50M is too large to be manually labeled, the dominant cluster detection quality can still be

qualitatively assessed by the results in Figure 10. As shown in Figures 10(b)-10(f), the affinity-based methods effectively detect most of the SIFT features (i.e., green points) extracted from the similar image regions of "KFC grandpa", since such SIFTs can naturally form dense subgraphs with high cohesiveness. At the same time, the large proportion of noisy SIFTs (i.e., red points) extracted from the random non-duplicate background image regions are filtered out. The results demonstrate the effectiveness of affinity-based methods in resisting overwhelming amount of noisy SIFTs. Additionally, we also evaluate the AVG-F performance of PALID on the labeled data sets of NART and NDI. The resulting AVG-F performances are consistent with ALID.

6. CONCLUSIONS

In this paper, we proposed ALID, a scalable and effective dominant cluster detection approach against high background noise. ALID demonstrates remarkable noise resistance capability, achieves significant scalability improvement over the other affinity-based methods and is highly parallelizable in MapReduce framework. As future work, we will further extend ALID towards the online version to efficiently process streaming data sources.

7. REFERENCES

[1] www.psi.toronto.edu/index.php?q=affinity%20propagation.

[2] https://sites.google.com/site/lhrbss/.

[3] P. Anderson, A. Thor, J. Benik, L. Raschid, and M. E. Vidal. Pang: Finding patterns in annotation graphs. In *SIGMOD*, pages 677–680, 2012.

[4] A. Angel, N. Sarkas, N. Koudas, and D. Srivastava. Dense subgraph maintenance under streaming edge weight updates for real-time story identification. In *PVLDB*, volume 5, pages 574–585, 2012.

[5] B. Bahmani, B. Moseley, A. Vattani, R. Kumar, and S. Vassilvitskii. Scalable k-means++. In *PVLDB*, volume 5, pages 622–633, 2012.

[6] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *JMLR*, 3:993–1022, 2003.

[7] J. Chen and Y. Saad. Dense subgraph extraction with application to community detection. *TKDE*, 24(7):1216–1230, 2012.

[8] W. Chen, Y. Song, H. Bai, C. Lin, and E. Y. Chang. Parallel spectral clustering in distributed systems. *TPAMI*, 33(3):568–586, 2011.

[9] L. Chu, S. Jiang, S. Wang, Y. Zhang, and Q. Huang. Robust spatial consistency graph model for partial duplicate image retrieval. *TMM*, 15(8):1–15, 2013.

[10] L. Chu, S. Wang, S. Liu, Q. Huang, and J. Pei. ALID: Scalable dominant cluster detection. in arxiv.org. 2014. Also available at the following URL. https://www.dropbox.com/s/nob6gsqas544yni/ALID%20Scalable%20Dominant%20Cluster%20Detection.pdf?dl=0.

[11] D. Comaniciu and P. Meer. Mean shift: A robust approach toward feature space analysis. *TPAMI*, 24(5):603–619, 2002.

[12] M. Datar, N. Immorlica, P. Indyk, and V. S. Mirrokni. Locality-sensitive hashing scheme based on p-stable distributions. In *SoCG*, pages 253–262, 2004.

[13] J. Dean and S. Ghemawat. Mapreduce: Simplified data processing on large clusters. *OSDI*, page 1, 2004.

[14] R. I. Dunbar. Neocortex size as a constraint on group size in primates. *Journal of Human Evolution*, 22(6):469–493, 1992.

[15] C. Fowlkes, S. Belongie, F. Chung, and J. Malik. Spectral grouping using the nystrom method. *TPAMI*, 26(2):214–225, 2004.

[16] B. J. Frey and D. Dueck. Clustering by passing messages between data points. *Science*, 315(5814):972–976, 2007.

[17] M. Li and J. T. Kwok. Making large-scale nystrom approximation possible. In *ICML*, pages 631–638, 2010.

[18] H. Liu, L. Latecki, and S. Yan. Fast detection of dense subgraphs with iterative shrinking and expansion. *TPAMI*, 35(9):2131–2142, 2013.

[19] H. Liu and S. Yan. Robust graph mode seeking by graph shift. In *ICML*, 2010.

[20] T. Liu, A. W. Moore, K. Yang, and A. G. Gray. An investigation of practical approximate nearest neighbor algorithms. In *NIPS*, pages 825–832, 2004.

[21] S. Lloyd. Least squares quantization in pcm. *Information Theory, IEEE Transactions on*, 28(2):129–137, 1982.

[22] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2):91–110, 2004.

[23] T. S. Motzkin and E. G. Straus. Maxima for graphs and a new proof of a theorem of turán. *Canadian Journal of Mathematics*, 17(4):533–540, 1965.

[24] A. Y. Ng, M. I. Jordan, Y. Weiss, et al. On spectral clustering: Analysis and an algorithm. In *NIPS*, pages 849–856, 2002.

[25] A. Oliva and A. Torralba. Modeling the shape of the scene: A holistic representation of the spatial envelope. *IJCV*, 42(3):145–175, 2001.

[26] M. Pavan and M. Pelillo. Dominant sets and pairwise clustering. *TPAMI*, 29(1):167–172, 2007.

[27] S. Rota Bulò and I. M. Bomze. Infection and immunization: a new class of evolutionary game dynamics. *Games and Economic Behavior*, 71(1):193–211, 2011.

[28] S. Rota Bulò, M. Pelillo, and I. M. Bomze. Graph-based quadratic optimization: A fast evolutionary approach. *CVIU*, 115(7):984–995, 2011.

[29] J. Sankaranarayanan, H. Samet, B. E. Teitler, M. D. Lieberman, and J. Sperling. Twitterstand: News in tweets. In *SIGSPATIAL*, pages 42–51, 2009.

[30] M. Shindler, A. Meyerson, and A. Wong. Fast and accurate k-means for large datasets. In *NIPS*, 2011.

[31] J. Sivic and A. Zisserman. Video google: A text retrieval approach to object matching in videos. In *ICCV*, pages 1470–1477, 2003.

[32] A. Vedaldi and B. Fulkerson. VLFeat: An open and portable library of computer vision algorithms, 2008.

[33] N. Wang, S. Parthasarathy, K.-L. Tan, and A. K. H. Tung. Csv: Visualizing and mining cohesive subgraphs. In *SIGMOD*, pages 445–458, 2008.

[34] F. L. Wauthier, N. Jojic, and M. I. Jordan. Active spectral clustering via iterative uncertainty reduction. In *SIGKDD*, pages 1339–1347, 2012.

[35] J. W. Weibull. *Evolutionary game theory*. The MIT Press, 1997.

[36] X. Xu, N. Yuruk, Z. Feng, and T. A. J. Schweiger. Scan: A structural clustering algorithm for networks. In *SIGKDD*, pages 824–833, 2007.

APPENDIX

A. PROOF OF PROPOSITION 1

PROPOSITION 1. *Given the local dense subgraph $\hat{x}$, the double-deck hyperball $H(D, R_{in}, R_{out})$ has the following properties:*

1. *$\forall j \in I$ and $||v_j - D||_p < R_{in}$, $\pi(s_j - \hat{x}, \hat{x}) > 0$; and*
2. *$\forall j \in I$ and $||v_j - D||_p > R_{out}$, $\pi(s_j - \hat{x}, \hat{x}) < 0$.*

PROOF. We refer to Equantion 15 in Section 4.2 as

$$\begin{cases} D = \sum_{i \in \alpha} v_i \hat{x}_i, \text{ where } v_i \in V \text{ are the data items.} \\ R_{in} = \frac{1}{k} \ln(\frac{\lambda_{in}}{\pi(\hat{x})}), \text{ where } \lambda_{in} = \sum_{i \in \alpha} \hat{x}_i e^{-k||v_i - D||_p} \\ R_{out} = \frac{1}{k} \ln(\frac{\lambda_{out}}{\pi(\hat{x})}), \text{ where } \lambda_{out} = \sum_{i \in \alpha} \hat{x}_i e^{k||v_i - D||_p} \end{cases} \tag{15}$$

Define $f_{in}(v_j)$ and $f_{out}(v_j)$ as

$$\begin{cases} f_{in}(v_j) = \lambda_{in} e^{-k||v_j - D||_p} \\ f_{out}(v_j) = \lambda_{out} e^{-k||v_j - D||_p} \end{cases} \tag{22}$$

By plugging λ_{in} and λ_{out} (Equation 15) into Equation 22, we have:

$$\begin{cases} f_{in}(v_j) = \sum_{i \in \alpha} \hat{x}_i e^{-k(||v_j - D||_p + ||v_i - D||_p)} \\ f_{out}(v_j) = \sum_{i \in \alpha} \hat{x}_i e^{-k(||v_j - D||_p - ||v_i - D||_p)} \end{cases} \tag{23}$$

Recall that the scaling factor $k > 0$ is always positive (Equation 1). Applying the triangle inequality to Equation 23, we obtain

$$\begin{cases} \forall j \in I, \pi(s_j, \hat{x}) \geq f_{in}(v_j) \\ \forall j \in I, \pi(s_j, \hat{x}) \leq f_{out}(v_j) \end{cases} \tag{24}$$

where $\pi(s_j, \hat{x}) = (A\hat{x})_j = \sum_{i \in \alpha} \hat{x}_i e^{-k||v_j - v_i||_p}$.

For any vertex v_j that satisfies $||v_j - D||_p = R_{in}$, we have

$$f_{in}(v_j) = \lambda_{in} e^{-kR_{in}} = \pi(\hat{x}) \tag{25}$$

by plugging R_{in} of Equation 15 into $f_{in}(v_j)$ of Equation 22.

Since $f_{in}(v_j)$ in Equation 22 monotonously decreases with respect to $||v_j - D||_p$, we can derive

$$\forall j \in I \text{ and } ||v_j - D||_p < R_{in}, f_{in}(v_j) > \pi(\hat{x}) \tag{26}$$

Therefore, we can derive from Equation 26 and the first inequation of Equation 24 that

$$\forall j \in I \text{ and } ||v_j - D||_p < R_{in}, \pi(s_j, \hat{x}) \geq f_{in}(v_j) > \pi(\hat{x}) \tag{27}$$

which proves the first property of Proposition 1.

The second property of Proposition 1 can be proved in a similary way by plugging $||v_j - D||_p = R_{out}$ of Equantion 15 into $f_{out}(v_j)$ of Equation 23. □

B. PROOF OF CONVERGENCE

PROPOSITION 2. *For a global dense subgraph x^* that contains M vertexes, denote $\hat{x}^{(c)}$ as the local dense subgraph detected by* Step 1 *of Algorithm 2 (i.e., ALID) in the c-th iteration, then we have the following two properties:*

1. *ALID is guaranteed to converge; and*
2. *The statistical expectation of the number of vertexes in $\hat{x}^{(c)}$ converges to M.*

PROOF. Some useful notations as listed as follow:

- x^* is the global dense subgraph containing M vertexes.
- $\hat{x}^{(c)}$ denotes the local dense subgraph detected by LID (i.e., Step 1) in the c-th iteration of Algorithm 2.
- $a^{(c)}$ denotes the statistical expectation of the number of vertexes in local dense subgraph $\hat{x}^{(c)}$.
- $H_c(D, R)$ is the ROI estimated from $\hat{x}^{(c)}$ in Step 2.
- $m^{(c)} \leq M$ represents the number of vertexes that belong to x^* and are within the ROI $H_c(D, R)$.
- $p \in (0, 1)$ is the lower bound of recall of the LSH method [12] in CIVS (i.e., Step 3). Theoretical proof of the lower bound of recall of LSH can be found in M. Datar's work [12].

First, we prove the convergence of ALID, which is the first property of Proposition 2.

According to Theorem 1 and Theorem 2, given any local range $\beta^{(c)}$, the LID method in Step 1 of Algorithm 2 is guaranteed to converge to the local dense subgraph $\hat{x}^{(c)} \in \triangle^n_{\beta^{(c)}}$ within the local range $\beta^{(c)}$. That is,

$$\hat{x}^{(c)} = \max_{x \in \triangle^n_{\beta^{(c)}}} \pi(x) \tag{28}$$

Recall that $\beta^{(c)}$ is updated by the vertexes retrieved by CIVS in Step 3 of Algorithm 2, we discuss how ALID converges in the following cases. For the interest of simplicity, we regard all the vertexes in the updated local range $\beta^{(c)} = a^{(c-1)} \cup \psi$ as retrieved by CIVS.

Case 1 : CIVS retrieves all the M vertexes of the global dense subgraph x^*. In this case, we have $x^* \in \triangle^n_{\beta^{(c)}}$. By the definition of global dense subgraph (See Equation 3), we can drive that

$$x^* = \max_{x \in \triangle^n} \pi(x) \tag{29}$$

Since $\triangle^n_{\beta^{(c)}} \subset \triangle^n$ and $x^* \in \triangle^n_{\beta^{(c)}}$, we have

$$x^* = \max_{x \in \triangle^n_{\beta^{(c)}}} \pi(x) \tag{30}$$

Therefore, we can derive from Equation 28 that LID is guaranteed by Theorem 1 and Theorem 2 to converge to the global dense subgraph $\hat{x}^{(c)} = x^*$. In other words, ALID converges to the global dense subgraph x^*.

Case 2 : CIVS retrieves N vertexes ($0 < N < M$) of the global dense subgraph x^*. In this case, the N retrieved vertexes of x^* form a subgraph $x' \in \triangle^n_{\beta^{(c)}}$, which is a subgraph of the global dense subgraph $x^* \in \triangle^n$. Considering the fact that all vertexes of x^* are highly similar with each other, we know that the vertexes in x' are highly similar with each other as well. Thus, the graph density $\pi(x')$ is the maximum in the local range $\beta^{(c)}$, that is

$$x' = \max_{x \in \triangle^n_{\beta^{(c)}}} \pi(x) \tag{31}$$

Therefore, we can derive from Equation 28 that LID is guaranteed to converge to $\hat{x}^{(c)} = x'$ in the c-th iteration of ALID. As proved later, the statistical expectation of the number of vertexes in $\hat{x}^{(c)}$ converges to M. In other words, ALID converges and the expected detection quality converges to the optimal result.

Table 3: The definitions of symbols and acronyms.

Symbols	Definitions	Acronyms	Definitions
n	The total number of graph vertexes.	ALID	The proposed method in this paper.
G	The global affinity graph $G = (V, I, A)$.	PALID	The parallel implementation of ALID
V	The set of all graph vertexes $V = \{v_i \in R^d \mid i \in I\}$. R^d is d-dimensional real number space.	LID	Localized infection immunization dynamics (Step 1 of Algorithm 2).
I	The "global range", which is the index set of all graph vertexes $I = [1, n]$.	ROI	Region of interest (Step 2 of Algorithm 2).
A	The global affinity matrix of G.	CIVS	Candidate infective vertex search (Step 3 of Algorithm 2).
v_i	The d-dimensional data item that uniquely corresponds to the i-th graph vertex.	LSH	The locality sensitive hashing method [12].
s_i	The n-dimensional index vector that uniquely represents the i-th graph vertex.	DS	The dominant set method [26].
x	The n-dimensional vector to represent any subgraph of G.	RD	The replicator dynamics method [35].
x_i	The i-th dimension of x. x_i is also the weight of vertex s_i in subgraph x.	StQP	Standard quadratic optimization problem (Equation 3).
$\gamma_\beta(x)$	The set of infective subgraphs that are infective against x and are within local range β. $\gamma_\beta(x) = \{y \in \triangle_\beta^n \mid \pi(y - x, x) > 0\}$.	IID	The infection immunization dynamics method [28].
a^*	The number of vertexes in the largest dense subgraph (i.e., dominant cluster).	SEA	The shrink and expansion algorithm [18].
α	The index set of all vertexes in subgraph x.	AP	The affinity propagation method [16].
β	The "local range", which is the index set of a local group of vertexes.	LSR	The locality sensitive region (Figure 4).
$A_{\beta\beta}$, $A_{\beta\alpha}$	Local affinity matrices (see Figure 3).	w.r.t.	This means "with respect to".
$c \leq C$	c is the current number of iterations of ALID. C is the maximum iteration limit of ALID.	AVG-F	Average F_1 score.
ω, η, P	The constant parameters in Table 1.	NDI	The near duplicate image data set.
δ	The maximum number of vertexes that can be retrieved by CIVS (Step 3).	Sub-NDI	The subset of NDI data set (Section 5.1).
$\triangle^n$, $\triangle_\beta^n$	$\triangle^n = \{x \in R^n \mid \sum_{i \in I} x_i = 1, x_i \geq 0\}$ is the set of all possible subgraphs in global range I. $\triangle_\beta^n = \{x \in \triangle^n \mid \sum_{i \in \beta} x_i = 1\}$ is the set of all possible subgraphs in local range β.	SIFT-50M	The data set of 50 million SIFT features [22].
$\hat{x}$, $\hat{x}^{(c)}$, x^*	$\hat{x}$ is the local dense subgraph found by LID. $\hat{x}^{(c)}$ is the dense subgraph detected by LID in the c-th iteration of ALID. x^* is the global dense subgraph (i.e., the output of ALID).	NART	The news articles data set.

Case 3 : CIVS retrieves zero vertex of the global dense subgraph x^*. This is an ill-conditioned case, which only happens when the recall of LSH [12] is $p \approx 0$ under improper parameters, such as extremely large number of hash functions and small segment length r (Figure 6). In this case, LID is still guaranteed to converge to a local dense subgraph $\hat{x}^{(c)} \in \triangle_{\beta(c)}^n$, which does not contain any vertex of x^*. However, such ill-condition can be easily avoided by setting the LSH parameters properly [12].

In summary, ALID is guaranteed to converge and the detection quality of ALID depends on how many vertexes of the global dense subgraph x^* is retrieved by CIVS.

In the following, we further prove that the statistical expectation of the number of vertexes in $\hat{x}^{(c)}$ converges to M, which is the second property of Proposition 2.

For Case 1, we have proved that $\hat{x}^{(c)} = x^*$, thus the second property of Proposition 2 is valid for Case 1.

For Case 2, we analyze the situation in detail as follows. By the definition of the global dense subgraph, all the M vertexes in the global dense subgraph x^* are highly similar with each other, therefore, when using a single graph vertex in the global dense subgraph x^* as the LSH query, the probability to retrieve each of the other vertexes in x^* is lower bounded by $p \in (0, 1)$ (i.e., the recall of LSH [12]).

Since CIVS uses all the $a^{(c)}$ different vertexes in the local dense subgraph $\hat{x}^{(c)}$ as LSH queries, we can derive that the probability to retrieve each of the vertexes in the global dense subgraph x^* is lower bounded by $1 - (1 - p)^{a^{(c)}}$.

Considering that there are totally $m^{(c)}$ vertexes of x^* inside the ROI $H_c(D, R)$, we can derive that the number of vertexes of x^* retrieved by CIVS is a random variable $b^{(c)}$ that submits to binomial distribution

$$b^{(c)} \sim B(m^{(c)}, 1 - (1 - p)^{a^{(c)}}) \tag{32}$$

Here, $b^{(c)}$ is the number of vertexes of x^* in the updated local range $\beta^{(c+1)} = \alpha^{(c)} \cup \psi$. As proved in Case 2, such vertexes form a subgraph $x' \in \triangle_{\beta^{(c+1)}}^n$ of the global dense

subgraph x^*, and LID is guaranteed to converge to the local dense subgraph $\hat{x}^{(c+1)} = x'$. Therefore, we can further derive from Equation 32 that the statistical expectation of the number of vertexes in $\hat{x}^{(c+1)}$ is $a^{(c+1)}$, where

$$a^{(c+1)} = E(b^{(c)}) = m^{(c)} \left[1 - (1-p)^{a^{(c)}}\right] \tag{33}$$

Recall that $m^{(c)}$ is the number of vertexes of x^* within ROI $H_c(D, R)$, since there are totally M vertexes in x^*, we have $m^{(c)} \leq M$.

For any iteration c, if $m^{(c)} < M$, then there are $(M - m^{(c)})$ vertexes of x^* outside the ROI $H_c(D, R)$ and we can derive

$$\hat{x}^{(c+1)} \neq x^* \tag{34}$$

from the fact that $\hat{x}^{(c+1)}$ is the local dense subgraph within the ROI $H_c(D, R)$ and $m^{(c)} < M$.

Then, we prove by contradiction that at least one of the $(M - m^{(c)})$ vertexes of x^* is infective against the local dense subgraph $\hat{x}^{(c+1)}$, as follows. If none of the $(M - m^{(c)})$ vertexes are infective against $\hat{x}^{(c+1)}$, then, as indicated by Theorem 1, $\hat{x}^{(c+1)}$ will be the global dense subgraph, which leads to $\hat{x}^{(c+1)} = x^*$. This is in contradiction with the fact that $\hat{x}^{(c+1)} \neq x^*$ in Equation 34. Therefore, at least one of the $(M - m^{(c)})$ vertexes is infective against $\hat{x}^{(c+1)}$.

Since the ROI $H_{(c+1)}(D, R)$ is guaranteed by Proposition 1 to obtain all the new infective vertexes against $\hat{x}^{(c+1)}$, we can derive that $m^{(c+1)} > m^{(c)}$. Therefore, we have $m^{(c)} < m^{(c+1)} \leq M$, which indicates that the series $\{m^{(c)}\}$ is an increasing series with a reachable upper bound M.

Recall that $p \in (0, 1)$, we can further derive from Equation 33 that the series of $\{a^{(c)}\}$ is an increasing series with respect to the iteration times c.

Since the total number of vertexes in the global dense subgraph x^* is M, we have $a^{(c)} \leq M$. Since $\{a^{(c)}\}$ is an increasing series, we can derive from Equation 33 that the series $\{a^{(c)}\}$ converges to M, and a larger value of p leads to a faster convergence rate. Recall that $a^{(c)}$ is the statistical expectation of the number of vertexes in local dense subgraph $\hat{x}^{(c)}$, thus, the second statement of Proposition 2 is proved. □

C. NOISE RESISTANCE ANALYSIS

The purpose of noise resistance analysis is to compare the noise resisting performance of the affinity-based methods (i.e., ALID, AP, IID, SEA) and the partitioning-based methods, such as: 1) k-means (KM) [21]; 2) spectral clustering using full affinity matrix (SC-FL) [24]; 3) spectral clustering using nystrom approximation (SC-NYS) [15]. For the partitioning-based methods, we use the source codes published by Chen *et al.* [8]. The mean-shift method (MS) [11] is analyzed as well.

We compare the AVG-F performance of each method on a series of data sets with increasing noise degree

$$\text{noise degree} = \frac{\#\ noise\_data}{\#\ ground\_truth\_data} \tag{35}$$

which is the ratio of the number of noise data items over the number of data items in the dominant clusters according to the ground truth. Such series of data sets are obtained by injecting a certain amount of noise data into the ground truth data set.

For partitioning-based methods KM, SC-FL and SC-NYS, we set the cluster number K in the same way as Liu *et al.* [18, 19] did, which counts the noise data as an extra cluster. We use the same settings as in Section 5 for the affinity-based methods and adopts the classic gaussian kernel for the mean shift method. All compared methods are optimally tuned. To avoid the performance degeneration caused by the enforced sparsity of the affinity graph, we use a full affinity matrix for methods AP, SEA and IID to preserve the original cohesiveness of all dense subgraphs. Due to the limited scalability of method AP, we use the two small data sets NART and Sub-NDI.

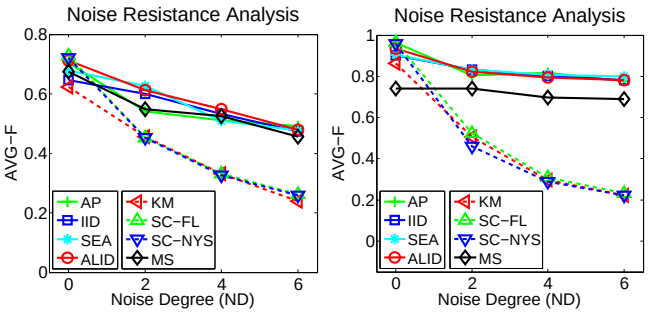

(a) NART (b) Sub-NDI

Figure 11: The noise resisting performances. (a) The results on the NART data set. (b) The results on the Sub-NDI data set.

As shown in Figure 11(a)-(b), the AVG-F performance of the partitioning-based methods decreases much faster than the other methods when noise degree increases. Such phenomenon is caused by the core mechanism of partitioning-based methods, which partitions all data items (including the noise data) into a fixed number of clusters. Thus, such mechanism is ineffective on noisy data. Similar observations is also reported by Pavan *et al.* [26]. In contrast, the affinity-based methods significantly outperform the partitioning-based methods on noisy data when the noise degree is large. This demonstrates the advantage of the affinity-based methods in resisting high background noise.

It is easy to understand why both SC-FL and SC-NYS achieve good AVG-F performance when "noise degree = 0", since those methods are fed with the correct number of clusters. However, the number of clusters is most likely unknown in practical applications.

As shown in Figure 11(a), MS achieves a comparable AVG-F performance with the affinity-based methods on the NART data set. This is achieved by a proper bandwidth setting of the pre-defined density distribution, which makes the MS method noise-resisting if the bandwidth properly fits the scale of most true clusters. However, on the Sub-NDI data set (Figure 11(b)), MS failed to achieve a comparable AVG-F with the affinity-based methods, since the distribution of the image features in Sub-NDI data set is more sophisticated than the text features in NART data set. Such performance degeneration is probably caused by the strong reliance of MS on the bandwidth and the pre-defined density distribution, which cannot simultaneously fit on various clusters under complex feature distributions.

In summary, ALID and the other affinity-based methods are more robust than the partitioning-based methods in

detecting unknown number of dominant clusters from noisy background.